P516765.PDF [Page 2 of 31]

DEFENCE R&D DÉFENSE

Mission-adaptive computer aided detection (CAD) for mine hunting: System concept and theory

Ronald T. Kessell

Defence R&D Canada

Technical Memorandum

DREA TM 2000-111

February 2001

National Defence Défense nationale

Canada

P516765.PDF [Page 1 of 31]

Mission-adaptive computer aided detection (CAD) for mine hunting: System concept and theory

Ronald T. Kessell

Defence Research Establishment Atlantic

Technical Memorandum
DREA TM 2000-111
February 2001

P516765.PDF [Page: 4 of 31]

Author

Ronald T. Kessel

Approved by

Dr. Ron Kuwahara
Head/EM Section

Approved for release by

Kirk Foster
Chair/Document Review Panel

© Her Majesty the Queen as represented by the Minister of National Defence, 2001

© Sa majesté la reine, représentée par le ministre de la Défense nationale, 2001

P516765.PDF [Page: 5 of 31]

104-07015/

Abstract

CAD systems require detection and rejection *training*—the adjustment of internal coefficients to discriminate between objects as required by the mission. Ideally, pattern recognition experts would carry out this training once, for all missions. But the range of mission parameters that affect detection performance is very large, making that ideal impossible to realize in practice. Final on-site adjustments are inevitable. This paper examines why this is the case, and then presents a concept for automatic on-site training, in which the CAD system gathers its training about targets and clutter at the start of a mission, non-intrusively, by "eavesdropping" on the sonar operator's routine responses to the sonar image (selecting potential threats and passing over clutter). This automated approach paves the way for more advanced CAD, whose pattern recognition schemes would otherwise require a prohibitive level of expert knowledge and effort to train them. It also improves the prospects for CAD by maintaining operational flexibility, and by focussing its range of application and training to just the mission at hand, rather than to all missions at once. CAD for sea mines using sidescan sonar is assumed here, but the concept applies to other detection problems as well.

Résumé

Les systèmes de détection assistée par ordinateur (DAO) exigent un apprentissage de la détection et du rejet — soit le réglage des coefficients internes pour distinguer les objets selon la mission. Idéalement, des experts en reconnaissance des formes effectueraient ce réglage une fois pour toutes les missions. Cependant, la gamme de paramètres de mission qui influent sur les performances en détection est très grande, ce qui rend l'idéal difficile à atteindre dans la pratique. Des réglages finals sur place sont inévitables. Le présent article examine pourquoi il en est ainsi et présente un concept pour un apprentissage automatique sur place, au cours duquel le système DAO acquiert son information au sujet des cibles et du fouillis en début de mission, de façon non intrusive, par ń écoute clandestine z des réponses courantes de l'opérateur sonar à l'image sonar (sélectionnant les menaces potentielles et laissant de côté le fouillis). Cette approche automatisée fraye la voie à une DAO plus perfectionnée, dont les combinaisons de reconnaissance des formes exigeraient autrement un niveau exorbitant de connaissance d'expert et d'effort d'apprentissage. Elle améliore également les perspectives pour la DAO en assurant la flexibilité opérationnelle et en concentrant sa gamme d'applications et d'apprentissage exclusivement sur la mission en cours, plutôt que sur toutes les missions à la fois. Il s'agit ici de la DAO de mines marines qui fait appel au sonar latéral, mais le concept s'applique également à d'autres formes de détection.

P516765.PDF [Page: 6 of 31]

Executive summary

Background

Activities in DRDC MCM are directed in part towards computer-aided detection (CAD) to help the sonar operator detect sea mines in sidescan sonar images. CAD entails pattern recognition, whose crucial first step is detection and rejection *training*—the adjustment of internal coefficients which enable the system to discriminate correctly between objects. In minehunting, one typically presents sonar images of known targets and clutter to the CAD system, while systematically adjusting its coefficients for optimal performance on that training set. Ideally this training would be done once, for all missions, by experts in pattern recognition. Unfortunately, the range of operational parameters that affect CAD performance is very large, making it unlikely that the performance-affecting conditions assumed for training will match those encountered during a mission. Some on site adjustments by the sonar operator are therefore required. These adjustments pose an added burden to the operator; a burden that is likely to increase as CAD algorithms become more complex in the bid for better performance.

Principal results

This paper presents the concept of CAD in which the final on-site mission adjustments are automatic and non-intrusive. The system gathers cues about targets and clutter by monitoring the actions of the sonar operator as they respond to the sonar image, much as they ordinarily would in the absence of CAD. By continually updating its training this way, the CAD system evolves towards the expert judgements of the operator, while signaling its state of readiness on the basis of past performance.

Significance of results

Automatic on-site training paves the way for more advanced, high performance CAD, without sacrificing mission flexibility or burdening the sonar operator. The mathematical relationship between performance and training, derived here to clarify and substantiate the relevant principles, should be of particular interest to CAD developers for sea minehunting and elsewhere.

Future Work

Preliminary tests of the concept will be carried out at DREA as candidate CAD algorithms become available.

Ronald T. Kessel. 2001. Mission-adaptive computer aided detection (CAD) for mine hunting: System concept and theory. DREA TM 2000-111. Defence Research Establishment Atlantic.

P516765.PDF [Page: 7 of 31]

Sommaire

Contexte

Les activités en LCM de RDDC portent en partie sur la détection assistée par ordinateur (DAO) pour aider l'opérateur sonar à détecter des mines marines sur des images de sonar latéral. La DAO implique la reconnaissance des formes, dont la première étape cruciale est l'apprentissage de la détection et du rejet — soit le réglage des coefficients internes permettant au système de distinguer correctement les objets. Dans la chasse aux mines, on présente normalement des images sonar de cibles et de fouillis connus au système DAO, tout en réglant systématiquement ses coefficients en vue de performances optimales pour ce jeu de conditions d'apprentissage. Idéalement, des experts en reconnaissance des formes effectueraient ce réglage une fois pour toutes les missions. Malheureusement, la gamme de paramètres de mission qui influent sur les performances en détection est très grande, de sorte qu'il est peu probable que les conditions influant sur les performances en DAO et utilisées pour l'apprentissage correspondent à celles rencontrées en mission. Par conséquent, l'opérateur sonar doit effectuer certains réglages sur place. Ces réglages augmentent la charge de travail de l'opérateur, charge qui accroîtra probablement à mesure que les algorithmes DAO deviennent plus complexes dans la poursuite de meilleures performances.

Principaux résultats

Cet article présente le concept d'une DAO dans laquelle les réglages finals sur place, en mission, sont effectués automatiquement et de façon non intrusive. Le système recueille des indices sur les cibles et le fouillis en surveillant les actions de l'opérateur sonar qui répond à l'image sonar, sensiblement de la même façon qu'il le ferait en l'absence de la DAO. En mettant à jour continuellement son apprentissage de cette manière, le système DAO évolue vers les jugements experts de l'opérateur, tout en signalant son état de préparation en fonction des performances antérieures.

Importance des résultats

L'apprentissage automatique sur place fraye la voie à une DAO perfectionnée à performances supérieures, sans sacrifier la flexibilité de mission ni surcharger l'opérateur sonar. La relation mathématique entre les performances et l'apprentissage, dérivée ici pour clarifier et corroborer les principes pertinents, devrait revêtir un intérêt particulier pour les développeurs de la DAO aux fins de la chasse aux mines et d'autres applications.

Travaux prévus

Des essais préliminaires du concept seront effectués au CRDA à mesure que des algorithmes DAO potentiels deviendront disponibles.

P516765.PDF [Page: 8 of 31]

Ronald T. Kessel. 2001. Détection assistée par ordinateur adaptative aux missions de chasse aux mines : Concept et principes du système . DREA TM 2000-111. Centre pour la Recherche de la Défence Atlantique.

P516765.PDF [Page: 9 of 31]

Table of contents

Abstract i

Résumé i

Executive summary ii

Sommaire iii

Table of contents v

List of figures vii

1. Introduction 1

1.1 The need for on-site adjustments 2

2. Training-Mission Mismatch 3

2.1 Operational factors that influence CAD 3

2.2 Route Survey, CF Kingston Class vessels 5

3. Improving CAD by refining its scope of application 5

3.1 Generalization: the relationship between performance and training 6

3.2 Optimal training 7

3.3 Ideal performance 7

3.4 Easy detection 8

3.5 Actual practice 8

4. Mission-adaptive concept for CAD training 9

4.1 A method for automated training 9

4.2 An illustrative example 10

4.3 Premises 13

4.4 Advantages 14

4.5 Disadvantages 15

4.5.1 Mitigation 15

P516765.PDF [Page: 10 of 31]

5. Research and development 16

6. Conclusions 16

References 17

P516765.PDF [Page: 11 of 31]

List of figures

1 Automatic Detection Training: The system acquires new or missed mine-like objects by "eavesdropping" on the sonar operator's response to the sonar image, and trains itself accordingly. The system monitors and reports its status based on past performance; the training continuing all the while regardless of the status "Not READY" would automatically switch to "READY" as training is acquired, when the probability of detection is high and false alarm rate is low. 11

2 Automatic Rejection Training: The system determines where its detections have erred by "eavesdropping" on the sonar operator's response to the sonar image, in this way verifying the systems alarms, identifying false alarms, and training itself accordingly. Rejection training is primarily for improved clutter rejection in response to CAD false alarms. The system monitors and reports its status based on past performance. 12

1. Introduction

When sidescan sonars are used for sea minehunting, computer aided detection (CAD) may be used to help the sonar operator scan the sonar image for mine-like objects; the goal being to improve the probability of detection by providing, in effect, a second pair of eyes whose attention never flags. All CAD systems use pattern recognition algorithms of one kind or another, and these algorithms must be "trained" before they are applied operationally—a process by which the internal decision boundaries, sensitivities, and thresholds of the pattern recognition scheme are adjusted to raise alarms for some classes of objects (detect mines), while passing over other classes (reject clutter). Thus "training" here refers to CAD system adjustment for optimal performance, not to human operator training.

In sea minehunting, a CAD system is usually trained by presenting sidescan sonar images (actual or computer generated) to the system, which include images of the mines sought and clutter to be rejected, and then by systematically tuning the system for optimal recognition performance on those training images. Ideally this procedure would be carried out once, for all missions, by experts in image processing and pattern recognition, so that CAD would be fully operational at the press of a button. That is how the best examples of pattern recognition have indeed worked in other applications, so it is natural to envision the same for CAD.

In practice, however, all existing CAD systems use some partial initial training by experts, but with further provisions made for final adjustments by the operator during each mission. A recent review of twelve existing CAD/CAC systems summarized matters by concluding that "it is reasonable to expect the sonar operator may tune the selected algorithm *on the order of, or less than, once per minute*, for the type of ocean bottom or ambient conditions encountered" (reference [1, Section 4.3.2], italics added here for emphasis). There is a discrepancy, then, between the ideal conception of CAD on the one hand, with its once-for-all approach to training, and the way it ultimately works in practice on the other hand, with its need for continual on-site adjustment.

This paper is the result of efforts to explain this discrepancy, and to discover the best approach towards improving CAD in light of it. A proper reckoning is required, for as it stands the continual adjustment of a CAD system calls for an operator who has working knowledge of the system's pattern recognition schemes, and who can adjustment them effectively, with little effort—a task, it is safe to say, which is likely to become more difficult as CAD algorithms are made more complex in the bid for better performance. Thus, many proposed advances to CAD algorithms risk doing more harm than good, giving worse net CAD performance rather than better, if they place ever greater demands on the operator for on-site training. This may in fact be the biggest hurdle to significant advances in CAD.

From a CAD developer's perspective, there are two ways around it: 1) to develop better algorithms that can be trained in the ideal once-for-all fashion, or 2) to automate the on-site adjustment process. In this paper it is argued that the wide range of operational

factors that can be encountered during minehunting make the prospects for the first option very poor; that this is, in effect, what the empirical evidence—the need for on-site training in all existing CAD systems—is telling us. The roadblocks to this first option are explored in detail here. They make the case for pursuing the second option: automated mission-adaptive training.

In this regard, military CAD systems typically enjoy one unique advantage over more conventional applications of pattern recognition, an advantage that CAD developers have perhaps viewed as a necessary concession to military rigor rather than as a potential benefit, and that is to have an expert operator on hand who is dedicated to performing the same recognition task as the system. With an operator at hand who is expert at finding mines in sonar images, it is possible for the CAD system to acquire new training and continually refine that training by watching over the operator's shoulder. That is, by non-intrusively verifying its own judgements against the operator's routine responses to the sonar image (selecting mine-like objects while overpassing clutter), in this way identifying its errors (missed targets and false alarms), and then activating its own on-site adjustments as necessary, with new corrective training instances drawn directly from the operator's console. This is a rough outline of the concept presented here.

Our approach is to show why on-site mission adjustments are necessary in practice, by first listing the dominant operational factors whose particular combination during a mission must be represented in the training sets (Section (2.1)), and then showing how the reduced scope of on-site training—reduced from many diverse missions to just the mission at hand—plays an enabling role for CAD (Section (3)). The concept proposed for automatic mission-adaptive training, including its essential algorithms, its premises, advantages, and disadvantages, is described (Section (4)). A way ahead for further development is also proposed (Section (5)).

1.1 The need for on-site adjustments

The reason why on-site adjustments have so far proved inevitable lies, we believe, in the balance that must be struck between operational diversity and performance. It is a balance treated more widely in pattern recognition under the term *generalization*—the application of rules conditioned by given instances (training) to classify new instances (performance). Generalization therefore refers to the relationship between performance and training, encompassing at once the relevance of the training to the application, the adequacy of the classifier, and the amount of information contained in the features used to discriminate between classes. And like any other application of pattern recognition, CAD fails to perform well when it fails to generalize, and it fails to generalize if

1. its training is not representative of its application, either because a) significant performance-affecting factors differ between the training set and operations in the field; or b) the training set is too small, being a random sample of targets and clutter, but not an adequate representation of their entire classes;

2. its classifier cannot handle the diversity of the training set, or

3. the features it extracts for discriminating between objects do not carry enough information[1].

It is shown below that effective on-site adjustment naturally follows the "path of least resistance" through all of these challenges, by minimizing the training-mission mismatch insofar as the training instances are now drawn from the mission itself (item 1a above); by using the most up-to-date training information as it becomes available during the mission (item 1b above); by reducing the scope of the application, and hence, the diversity of the training, to just what is required for a particular mission, which in turn plays an enabling role for the classifier (item 2 above); and by constraining the relevant features , reducing their spread and overlap between the target and clutter classes[2] (item 3 above). CAD developers have in fact followed this path whenever they have made provisions for on-site adjustments by the operator, either as an intentional design feature from the outset, or as a concession to the difficulties encountered when the system was tested operationally.

2. Training-Mission Mismatch

As with most Naval systems, CAD for minehunting faces a very wide range of operational factors, many of which can affect the quality of the target signatures and background clutter. An experienced human operator can account for these factors very quickly, almost at a glance, whereas a CAD system always judges on the basis of its previous training, carrying the rigid bias of that training into the mission. To achieve the high levels of performance required of CAD, the operational factors assumed in the training must closely match those encountered in the mission.

2.1 Operational factors that influence CAD

Many likely causes of training-mission mismatch with sidescan sonar are listed below.

1. Type of mines: There are of course many types of mines. They can roughly be classified as bottom, tethered, or floating, and further by their shape, cylindrical, conical, spherical, or stealth. Training must include all of the types expected on a mission.

2. Make and model of the sonar: All sidescan sonars are not the same. Image quality, especially highlight-shadow contrast and pixel resolution, can differ dramatically

[1] In the language of information theory, the entropy of the features may be too high, resulting in unsatisfactory detection performance no matter how good the classifier or training may be. New discriminatory features must be found in such cases. Failing that, the range of application might be reduced in scope from many diverse scenarios to just one, with lower overall entropy resulting if the features are then more tightly clustered, with better class separation

[2] Decreasing the entropy of the feature set.

between sonar manufacturers, and can even differ appreciably between sonars of the same make and model owing to variance in transducer sensitivities. Training images for CAD should therefore, at the very least, be collected by sonars of the same make and model as those used operationally. Ideally, the same sonar would be used for both training and operations, which is in fact achieved in the proposed CAD training concept.

3. Towfish flight controls: A sidescan sonar is towed under water behind a ship at a speed of about 10 knots. Some towfish have dynamic flight controls for terrain-following mode, and limiting pitch and yaw. Others have much weaker passive flight control, through the aerodynamic design of a depressor wing for instance. Dynamic towfish control typically results in better sonar stability, and hence, better images. Training images should therefore be collected using the same towfish control as that used operationally.

4. Sonar altitude: The altitude of the sonar affects the viewing angle of targets, and hence their appearance in the sidescan sonar image. The same altitude should therefore be used for collecting training images as will be used operationally. This may be very difficult when dynamic flight controls are not used, especially when training is carried out shallow water and operations are carried out in deep, or visa versa.

5. Sonar's adaptive display sensitivity: Many sonar systems adjust the absolute sensitivity of the visual display such that the average image intensity is uniform across range and along track, and is consistently bright whether the seafloor is low reflective or high. This is called histogram equalization or transformation for improved visual contrast. The apparent intensities of the signature of a target can therefore change somewhat, depending on the backscatter strength of the seafloor, with high backscatter leading to suppressed target echoes, and low backscatter leading to enhanced, the change being purely the result of background dependent image processing rather than to acoustic sensitivity. The CAD system must nevertheless take this adaptive sensitivity into account, either by processing pre-visual sonar data before any adaptive histogram transform is applied, or, when only the visual image is available, somehow adjusting its training adoptively.

6. Sonar resolution setting: Quite apart from the question of make and model (item 1 above), sidescan sonars typically have two or more resolution settings to permit a fast low-resolution search for preliminary surveys, or a high-resolution search for hunting. Training should therefore be carried out with the same resolution setting as that used operationally. Otherwise CAD must be made resolution invariant.

7. Environmental parameters: Optimal detection performance cannot be achieved by focussing on the target class alone; it must also include clutter rejection [6]. But local geology, sedimentation, and biological life (plants and fish) can affect background clutter dramatically, and (to a lesser extent) so can water salinity and temperature, making it very difficult to train for clutter rejection in advance of a mission. This is true for Canadian waters, and it is especially true for forward

operations outside Canadian waters. Strong ocean currents can also lead to inordinately high or low above-ground survey rates, which could affect image resolution, and lead to increased turbulent sonar motion and a corresponding image degradation.

2.2 Route Survey, CF Kingston Class vessels

It should be noted that the Route Survey systems now operating on the Canadian Kingston Class vessels go a long way towards operational uniformity so far as the factors 2 to 6 above are concerned. Thus all route survey systems are equipped with sonars of the same make and model, and all have equal dynamic towfish control to keep the sonar locked at predefined altitudes in terrain-following mode.

Questions about flexible CAD training nevertheless arise, mainly because the Kingston Class was not designed for passage through mined waters. A technology demonstrator project now underway, the Remote Mine-Hunting System (RMS), is a promising alternative for instance. It remains uncertain, then, whether a different sonar (more expendable perhaps) will be used for minehunting as that used for route survey; or, if the same sonar is used, whether the same dynamic flight controls and sonar altitudes can be achieved.

The Route Survey system also collects image "mug shots" of mine-like objects into a database maintained by the Route Survey Office, and these could be used for CAD training. But the images are presently stored after geocoding, that is, position referenced and rotated to "north-up". This is required for route survey, but presents complications for CAD more generally inasmuch as it requires pattern recognition schemes that are 360-degree rotationally invariant, even when targets are axially symmetric.

3. Improving CAD by refining its scope of application

If we nevertheless assume for the moment that the training set is fully representative of all missions, then CAD still faces the challenge of performing well given such a wide range of operational scenarios. The complexity of the classifier necessarily increases as the diversity of the scenarios increases, and so do the demands placed on the features that CAD uses to discriminate mines from clutter. Thus the prospects for successful CAD improve when the scope of its training is reduced to just what is required for the mission at hand. This is the advantage of a specialist focussed on a narrowly defined task (mission-adaptive training), over a generalist with a much wider scope of attention (once-for-all training). It takes an exceptional generalist to outperform a specialist in a challenging task.

The principle is demonstrated here from a CAD developer's perspective using a mathematical treatment. A probabilistic approach is required because the

characteristics of new instances encountered by CAD cannot be predicted with certainty. The analysis applies generally to the training of any CAD system, whenever its internal coefficients are adjusted for optimal performance in some way, no matter how simple or complex its pattern recognition algorithms may be. It also serves as a basis for more detailed analyses of particular CAD systems.

3.1 Generalization: the relationship between performance and training

Assume first of all that the training is to prepare the CAD system for M identifiable operational scenarios, labeled $m = 1 \ldots M$, each being a particular realization of the performance-affecting factors listed in Section (2.1) for instance. This list depends in part on the type of training we choose. In a once-for-all approach to training, it includes all of the likely scenarios that might be encountered throughout the life of the CAD system, whereas, in mission-adaptive training, it only includes the scenarios that are likely to be encountered during the particular mission at hand. In either case, let q_m be the relative frequency with which the m'th scenario is encountered. These may also be interpreted as prior probabilities in what follows.

Let τ_n, $n = 1 \ldots N$, be N internal CAD system parameters that are to be adjusted by training, and let $P = P(\tau_1, \tau_2, ..., \tau_N)$ be a measure of the performance by which the CAD system is judged, often the probability of detection, though not necessarily[3]. The expected performance for a given training is the mean

$$\langle P \rangle = \sum_{m=1}^{M} q_m p_m, \tag{1}$$

in which $p_m = p_m(\tau_1, \tau_2, ..., \tau_N)$ is the performance given the m'th scenario alone.

Equation (1) shows how the overall performance of the CAD system depends on the prevalence of each scenario (the relative frequencies or prior probabilities q_m), and on the ability of the CAD system to train for each scenario (by optimizing each p_m). It therefore represents the process of generalization inasmuch as the performances p_m at one operational level, determined by the training process, are taken to a higher mission level comprised of the lower in relative proportions q_m.

Part of the difficulty in the analysis of training stems from the fact that this kind of generalization actually takes place at many different levels at once. The individual performance measures p_m could themselves be written in a form (1) with the summation now over a series of more narrowly defined "sub-scenarios", characterized

[3] P could be the probability of false alarm to be minimized, for instance, or some measure of the operational cost or risk, likewise to be minimized. It is usual, moreover, to speak of optimal performance given certain constraints, such as the maximization of the probability of detection for a given probability of false alarm, as under the Neyman-Pearson criteria. The optimization constraints derived later in Section (3.2) for the single measure P can be extended to multiple performance constraints using the method of Lagrange multipliers.

now at lower operational levels, such as target range and viewing aspect, which do not distinguish the original scenarios but nevertheless play a determining role in the performance for each. Those sub-scenarios might in turn be subdivided into more narrowly defined "sub-sub-scenarios", and so forth. But at some stage we must say that the training set is representative of an identifiable grouping of performance-affecting factors, which we call the scenarios, and for which we can estimate their individual performances p_m through the training process.

Much can be done with equation (1) to demonstrate important principles of generalization, such as the need for the training set to be representative of the target population rather than a mere random sampling of it (an instance of pitfall 1b, pg. 2), and the phenomenon of *over training* when a classifier focuses too particularly on individual training instances and fails to generalize beyond them (pitfall 1a). Here we demonstrate the difficulties that training faces in mathematical terms, in order to show how mission-adaptive training avoids them as much as possible, whereas once-for-all training confronts them head on.

3.2 Optimal training

The goal of training is to find coefficients τ_n which optimize the expected performance $\langle P \rangle$. Using variational methods in optimization, $\langle P \rangle$ is optimized with respect to τ_n when

$$0 = \frac{\partial \langle P \rangle}{\partial \tau_n} = \sum_{m=1}^{M} q_m \frac{\partial p_m}{\partial \tau_n}. \tag{2}$$

This must be true for all n, giving a system of N constraints that strike an optimal balance between the constituent scenarios; the performance of unlikely scenarios being sacrificed for better performance on more likely scenarios.

Needless to say, this optimal training ignores the possibility of training-mission mismatch considered above in Section (2). In mathematical terms, such mismatch implies that the scenario list $m = 1 \ldots M$ is incomplete, and the mission is not represented among the training constraints (2).

So far as the classifier is concerned, its greatest difficulty is that of satisfying the N in (2) constraints simultaneously. As a general rule, the more diverse the scenarios are, the more independent the constraints become, and the more complex the classifier must be to achieve the optimal performance.

3.3 Ideal performance

Many approaches to CAD training avoid mentioning prior probabilities altogether by aiming for ideal performance from the outset; that is, optimized performance for *any*

P516765.PDF [Page: 19 of 31]

combination of scenarios. The constraints (2) can be made insensitive to q_m by optimizing the individual performances p_m independently, such that

$$0 = \frac{\partial p_m}{\partial \tau_n}, \tag{3}$$

for all m and n. The original N constraints (2) are obviously satisfied for all q_m when (3) is true. Thus it is evident that ideal generalization to all missions requires that the classifier used by the CAD system has sufficient complexity that it is capable of optimizing each operational scenario independently (which amounts to satistying $M \times N$ constraints simultaneously), and that the training set be fully representative of operations in the field (i.e., the training scenario list $m = 1 \ldots M$ is complete). Here again, the more diverse the scenarios are, the more independent the constraints become, and the more complex the classifier must be to achieve the ideal performance.

It is important to note, that, although the scenario prior probabilities can be ignored for ideal training, they remain inescapable in a larger operational context, reappearing on the one hand when the performance (1) resulting from the ideal training is estimated when planning a mission, and on the other hand when the performance p_m for a particular scenario is further analyzed in terms of its constituent "sub-scenarios" as mentioned earlier. The outright omission of prior probabilities at every lower operational level or sub-scenario ultimately raises single training instances to the level of scenarios, resulting in over training, such as one must guard against when training neural networks and decision trees for example.

3.4 Easy detection

Ideal performance can be approximated when the target and clutter classes are well separated in feature space. Good separation depends in part on the features chosen for discrimination between target and clutter classes. Thus an easy detection problem becomes hard when the wrong features are used.

Good separation implies that the internal CAD coefficients τ_n have a wide range of values over which the individual scenario performances p_m are high and do not change significantly, and hence that the $M \times N$ constraints (3) can be simultaneously satisfied with little effort, almost as a matter of course. There are realistic scenarios where this is true in minehunting, such as for bottom mines on a hard uncluttered seafloor, in which case detection by a human operator's visual inspection of the sonar image is likewise very good.

3.5 Actual practice

Researchers naturally set ideal performance as their goal, and they begin with easy detection problems—data sets in which the targets are visually obvious against a benign background. But what one typically finds in practice is that a CAD system trained for optimal performance for one such scenario fails when applied to another. And when the

training is adjusted for that second scenario, it fails again for the first. Indeed, it is not unusual for the researcher to readjust the training for virtually every scenario tested, much as the sonar operator has so far been required to make in the existing CAD systems. The trouble here is that the CAD system does not generalize, not even from one part of an experimental training set to another.

The three causes for such failures were listed in Section (1, page 2). The first, mismatch between the training and operations, can be ruled out for the moment because the training and operational test sets are one and the same at the early stage in the development and testing of CAD considered here. The second and third remain: that the classifier is too simple and cannot satisfy all constraints (3) simultaneously, or that the features used for discrimination are inadequate for the desired performance $\langle P \rangle$ (i.e., they make detection insurmountably difficult). Continual mission-adaptive training alleviates these two difficulties as much as possible by focusing the training and application of CAD to just the mission at hand. If we are to judge by the existing CAD systems, moreover, the effect of continual on-site adjustment on detection performance can be significant, raising poor performance levels to acceptable levels. But the ability to make effective on-site adjustments, continually throught a mission, requires considerable skill and experience on the part of the operator, acquired as much by trial and error as by analysis it would seem.

4. Mission-adaptive concept for CAD training

Granted that on-site adjustments are inevitable, then, we now consider how they can be automated for maximum effectiveness.

4.1 A method for automated training

One can imagine the new CAD system starting its first mission with some nominal initial pattern recognition training, but without the on-site adjustments needed for optimal performance for the given mission. The mission begins, and the operator watches the scrolling sonar image for mine-like objects, marking potential mine threats for further consideration much as they ordinarily would in the absence of CAD.

Meanwhile, the CAD system monitors the operator's on-screen selections, to learn what kind of objects are of particular interest (detection training), and what objects are left unselected, to learn what kinds of clutter should be rejected (rejection training). Mission-adaptive training is acquired this way, with no action from the operator beyond what would in any case be expected of them during the search.

Algorithms for the detection and rejection training are given Figs.1 and 2, respectively. When the operator marks the first mine-like object, the CAD system automatically acquires the type of targets that are of interest to the operator, in the manner shown in Fig. 1, and the CAD system can then begin its task of helping the operator to detect

those objects, by automatically marking potential targets on the scrolling sonar image for instance.

Training continues as the CAD system continues to monitor the operator's actions, comparing its judgements with the operator's, learning both to acquire new classes of targets as they are encountered, and to reject clutter as the operator de-selects—perhaps with a single on-screen mouse click—false alarms raised by the CAD system that are, to the operator's mind, obviously false in view of the mission objective and its operational parameters. This rejection training process is shown schematically in Fig. 2.

In this way the CAD system constantly improves its training, learning to detect targets and minimize false alarms as the mission progresses. At no time is the operator required to enter specialized image processing or pattern recognition parameters.

4.2 An illustrative example

More than one existing CAD system relies on *energy detection*—a simple detection scheme that registers the presence of a mine-like target whenever the echo strength exceeds the mean backscatter level by a prescribed amount, and for a given duration (i.e., a given number of across-track pixels in the sidescan sonar image). This is a simplistic approach to detection, but it serves to illustrate the need for, and method followed in, the present automated concept.

As it stands with such a system, the sonar operator must manually enter the threshold echo level and duration, which, when exceeded, will trigger the detection alarm. And he must enter them repeatedly during a mission because the relative importance of the mean backscatter strength and echo highlights are continually changing in response to variable seafloor properties, background clutter, and so forth. In terms of the preceding analysis, continual adjustment is required because both the classifier and the feature set are inadequate for "hands-free" operation of the system based on a single training session. The method can nevertheless be shown to perform adequately on many occasions, provided that the thresholds have been carefully adjusted. The operator's attention is therefore split between the primary task of scanning the image for mine-like objects, and estimating whether or not CAD adjustments are required, and then making those adjustments in an effective way.

Now imagine such a system operating under a simplified version of the automated concept proposed here. Rather than entering thresholds that require a working knowledge of the CAD system, the operator concentrates fully on the task of recognizing mines in the sonar image, which he registers by clicking a mouse-driven cursor over the object's signature in the image. For simplicity in this example, let the CAD system begin the mission in a "Non-Ready" state, merely waiting for the operator to select the first prospective target, making no efforts towards its own rejection training. In time, the operator recognizes a potential threat, and the action of selecting it on the screen with the mouse not only flags the contact for further operational

P516765.PDF [Page: 22 of 31]

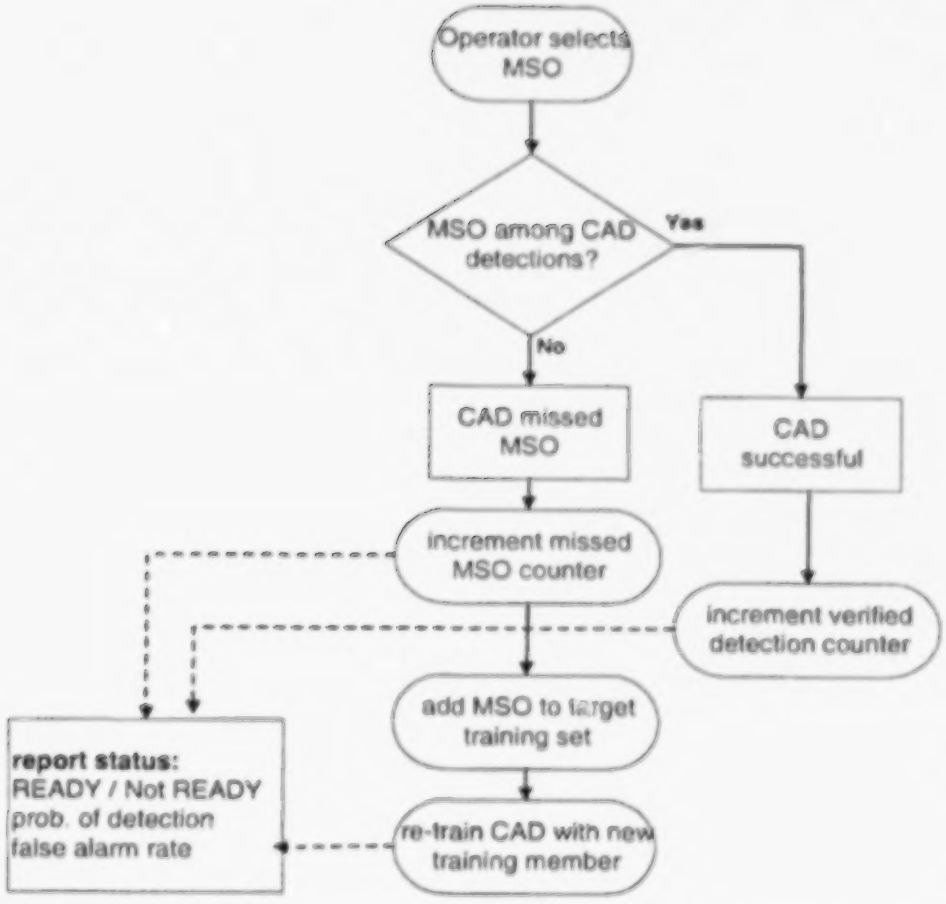

Figure 1: *Automatic Detection Training: The system acquires new or missed mine-like objects by "eavesdropping" on the sonar operator's response to the sonar image, and trains itself accordingly. The system monitors and reports its status based on past performance, the training continuing all the while regardless of the status. "Not READY" would automatically switch to "READY" as training is acquired, when the probability of detection is high and false alarm rate is low.*

P516765.PDF [Page: 23 of 31]

***Figure 2:** Automatic Rejection Training: The system determines where its detections have erred by "eavesdropping" on the sonar operator's response to the sonar image, in this way verifying the systems alarms, identifying false alarms, and training itself accordingly. Rejection training is primarily for improved clutter rejection in response to CAD false alarms. The system monitors and reports its status based on past performance*

consideration, but it also signals the CAD system that here is a target from which it can acquire its first detection training. The CAD system then analyzes the selected object to automatically determine—by exhaustive trial-and-error search of all possible threshold levels if necessary—what thresholds are required for the CAD system to register the same detection, on the basis of the energy detection principles alone, without raising too many false alarms. Having acquired its training by non-intrusively "eavesdropping" on the sonar operator this way, the CAD system switches to a "Ready" state, alerting the operator that it is ready for duty.

This simplified version of the automated training concept may in fact go a long way towards improving the performance of the energy detection scheme. But as we have seen, training is generally more complex, requiring more than a single instance for detection training, and requiring rejection training at the same time. This is particularly true of more advanced detection schemes. The concept algorithms in Figs. 1 and 2 address these more complex operations. If applied to the simple energy detection scheme here, the CAD system would automatically find the threshold settings that replicate the judgements of the operator about targets and clutter as much as possible, given the limitations of the energy detector's simplicity. Thus they optimize performance, without requiring the operator to enter any specialized parameters of pattern recognition. These are the two main goals of the automated concept.

Although the CAD system aligns its judgements about targets and clutter with the operators, the new alarms it raises while in its "Ready" state are based on a very thorough, independent assessment of every pixel in the sidescan sonar image, serving as a second pair of eyes for the operator, drawing attention to threats in portions of the image that were perhaps overlooked, or where the clutter was visually overwhelming, in this way improving the probability of detection.

4.3 Premises

The concept assumes that:

1. the operational parameters relevant to pattern recognition in minehunting are too many and too uncertain to be covered by a single once-for-all approach to training-training must instead proceed on a mission-to-mission basis;

2. the human operator's judgement is the ultimate standard of performance for CAD to achieve; in other words, that the CAD system should perform with roughly the same discrimination abilities as the expert human operator, serving as a second pair of eyes whose attention is thorough, never distracted, and never tires;

3. it is technologically feasible for CAD systems to acquire their training quickly, in real time, with relatively small training sets (small compared with pattern recognition generally).

Of the second premise it should be said that continual real time training has the best

P516765.PDF [Page: 25 of 31]

prospect of reaching the goal of roughly equal discrimination performance as a human operator. Of the third it should be said that fast training with small training sets becomes more feasible as the scope of the training is reduced from many diverse scenarios to a single mission, as described in Section (3).

4.4 Advantages

The CAD system would:

1. adapt itself automatically to a wide range of minehunting operations (different sonar manufacturers, sonar settings, environments, targets of interest, and others listed in Section 3 below), without the need for a major CAD retraining effort, giving maximal flexibility for minehunting operations,
2. free the operator from the burden of adjusting numerical parameters that require a working knowledge of image processing and pattern recognition (intensity thresholds, pixel counts, and so forth);
3. work as much as possible towards achieving the same ability of discrimination as the human operator, so far as the difficulty of each particular mission allows,
4. easily acquire detection training for new or unexpected mine types (bottom or tethered for instance);
5. easily acquire detection training for other objects of interest, such as a pipelines, lost cargo, or wreckage;
6. accommodate many different types of pattern recognition: thresholding, highlight-shadow matching, correlation filters, neural networks, etc.; while also making it possible to use more advanced algorithms that would otherwise burden the operator with specialized on-site adjustments;
7. be self optimizing to some degree, by always adjusting itself towards achieving a performance equal to the operator's,
8. report its performance in terms of empirically derived probabilities of detection and false alarm rates ("high", "low", or numerical measures), as well as overall CAD readiness ("Ready" or "Not Ready"), computed on the basis of performance to date by comparing its own judgements against the actions of the human operator;
9. continue learning and updating CAD status reports by monitoring the actions of the operator, even when CAD on-screen alerts are switched off for a time, for instance, when the false alarm rate may be distractingly high at the very start of training;
10. automatically collect image training libraries under the expert supervision of sonar operators, which are guaranteed to match at subset of realistic Naval operations rather than staged experiments, and which may therefore serve in further development of Computer Aided Classification (CAC) and future autonomous minehunting systems;

P516765.PDF [Page: 26 of 31]

11. function independently from a centralized CAD training centre by acquiring its training on site, from scratch if necessary, or borrowing shared initial training acquired by other systems used on similar missions;

12. spare Operations the considerable task of collecting and maintaining diverse training libraries before CAD can be made operationally useful.

4.5 Disadvantages

On-site training of any kind, whether by the human operator or automated, suffers from some delay from the time that new training begins until it is fully acquired. Automatic training may therefore lapse into a "non-ready" state, at the start of a mission or as radically different clutter is encountered, while its false alarm rate may be unacceptably high and targets may be missed. Similar lapses (or worse) would of course occur with more conventional CAD systems, whenever the operator is obliged to adjust the training coefficients of the CAD system manually, likewise at the start of a mission or background clutter changes, by trial and error as often happens. The automatic approach may therefore be the most effective approach to such lapses from the operator's perspective.

4.5.1 Mitigation

Here again, the reduced scope of application, from many diverse scenarios to just the mission at hand, should enable faster and more effective training than one could expect when training for all diversity at once, because the demands on the classifier and features are reduced (Section (3)) .

At the same time, one could use pattern recognition algorithms that accommodate diverse background clutter (sand, rocks, etc.) by discriminating between more than just two classes (target/non-target), such as are required in any case—whether automated or not—when more than one mine type is of interest (bottom and tethered mines, for instance). In effect, multi-class capabilities of this kind would serve as a kind of memory of targets and clutter types encountered during a mission, which would be ready for immediate use whenever those types are encountered again. Here again, limiting the scope of mission-adaptive training to a single mission keeps the number of target and clutter classes to an absolute minimum, making multi-class pattern recognition as simple as possible.

In the present concept, then, lapses can be minimized by striking a balance between the speed of training and the added complexity of multi-class discrimination. At the same time, the system's performance measures could alert the operator when a lapse occurs.

P516765.PDF [Page: 27 of 31]

5. Research and development

The concept requires a pattern recognition scheme that can be trained in real time with small training sets. Schemes of this type are in use for minehunting [2, for example] and SAR imagery [3, for example], and others have received preliminary testing by the MCM Group at DREA [4][5][6]. Tests are underway, moreover, with a commercially available neural network that shows promise.

Successful pattern recognition depends to a large extent on the features used for discrimination. Research continues at DREA into the best features that can be extracted from sonar images for the detection and classification of mine-like objects. These include non-linear pre-processing; wavelet and other transforms; shadow-highlight size, shape, intensities, contrast; and visual cues used by human operators.

Preliminary functional tests of the concept could be carried out with actual sonar images, using Windows NT personal computers together with image processing and high level programming tools such as IDL, ENVI, Matlab, Mathcad, LabView, and HNet2000.

Topics that require further investigation in support of the concept are:

1. multi-class detection and rejection when several very different classes of targets (bottom mines and tethered mines for instance) are sought in the presence or different clutter types (rock or sand, for example). [7][8][9];

2. unbiased statistical estimates of performance, and their optimization through automatic mission-adaptive training [6].

6. Conclusions

Unlike many applications of pattern recognition, CAD in minehunting faces unique challenges that make it very difficult, if not impossible, to achieve acceptable performance for all missions with just one, or an occasional, training adjustment to the system. CAD developers have of necessity turned to on-site mission adjustments to minimize these difficulties as much as possible. On-site adjustments must be kept as simple and straightforward possible from a sonar operator's perspective, or else they are likely to pose an operational burden that dramatically reduces the net benefit of the CAD system.

A concept for automatic mission-adaptive training was proposed here, which exploits the one advantage that CAD has above most other applications of pattern recognition, namely, to have a human operator on hand who is dedicated to the same minehunting task. Thus, by non-intrusively "eavesdropping" on the operator's routine responses to the sonar image (selecting targets and passing over clutter), the CAD system can judge

P516765.PDF [Page: 28 of 31]

its own performance, and initiate its own corrective adjustments based on instances drawn from the operator's console, viewing the operator as a teacher from whom it gathers its training.

This automated approach provides the benefits of on-site training—minimum training-mission mismatch, with maximum prospects for high performance—without burdening the sonar operator. Thus it paves the way towards advanced CAD algorithms whose complexity would otherwise make on-site training prohibitive.

References

1. Wort, P., CAD/CAC Study Report, MacDonald Dettwiler Report CL-RP-50-7394, Issue 1, Revision 0, June 27 1996.

2. Dobeck G.J., Hyland J.C. , and Smedly L., "Automated detection/classification of sea mines in sonar imagery," Naval Research Reviews, Office of Naval Research, V.49, pp. 9-20, 1997.

3. Shenoy R., and Casasent D., "Correlation filters that generalize well," SPIE 3386, pp. 100-110, 1998.

4. Fawcett J. , "Preliminary results of an algorithm for automatic detection of mine-like objects in sidescan sonar images", Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 2000-013, Jan. 2000.

5. Kessel R.T. , "Trainable linear filters for the detection of mines in side scan sonar images," Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 1999-171, Dec. 1999.

6. Kessel R.T. , "A probabalistic theory for the design of optimal linear discriminators for the automated detection of objects in sidescan sonar images," Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 2000-021, Feb. 2000.

7. Fukunaga K., "Introduction to statistical pattern recognition", 2nd Ed., Academic Press, New York, 1990.

8. Devroye L. , Gyorifi L. , and Lugosi G., "A probabilistic approach to pattern recognition", Springer, New York, 1996.

9. Mahalanobis, A., Vijaya Kumar B.V.K., and Sims S.R.F., "Distance classifier correlation filters using multiple eigenvectors for enhanced class separation", SPIE 2565 pp. 50-61, 1995.

P516765.PDF [Page: 29 of 31]

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1 ORIGINATOR (the name and address of the organization preparing the document Organizations for whom the document was prepared, e g Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8)

Defence Research Establishment Atlantic
PO Box 1012
Dartmouth, Nova Scotia, Canada B2Y 3Z7

2 SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

UNCLASSIFIED

3 TITLE (the complete document title as indicated on the title page Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title)

Mission adaptive computer aided detection (CAD) for mine hunting System concept

4 AUTHORS (Last name, first name, middle initial If military, show rank, e g Doe, Maj John E)

Kessel, Ronald T.

5 DATE OF PUBLICATION (month and year of publication of document)	6a NO OF PAGES (total containing information Include Annexes, Appendices, etc)	6b NO OF REFS (total cited in document)
Feb 2001	23 (approx.)	9

7 DESCRIPTIVE NOTES (the category of the document, e g technical report, technical note or memorandum If appropriate, enter the type of report, e g interim, progress, summary, annual or final Give the inclusive dates when a specific reporting period is covered)

TECHNICAL MEMORANDUM

8 SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development Include address)

Defence Research Establishment Atlantic, Electromagnetics Section, MCM Group

9a PROJECT OR GRANT NO (if appropriate, the applicable research and development project or grant number under which the document was written Please specify whether project or grant)	9b CONTRACT NO (if appropriate, the applicable number under which the document was written)

10a ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity Thus number must be unique to this document)	10b OTHER DOCUMENT NOs (Any other numbers which may be assigned this document either by the originator or by the sponsor)
DREA TM 2000-111	

11 DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12 DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document This will normally correspond to the Document Availability (11) However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected)

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

13. ABSTRACT (a brief and factual summary of the document It may also appear elsewhere in the body of the document itself It is highly desirable that the abstract of classified documents be unclassified Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U) It is not necessary to include here abstracts in both official languages unless the text is bilingual)

(U) CAD systems require detection and rejection training—the adjustment of internal coefficients to discriminate between objects as required. Ideally this machine training would be carried out once for all missions by pattern recognition experts. But the range of mission parameters that affect CAD training is very large, making a once-for-all approach particularly difficult, if not impossible. Final on-site mission adjustments are inevitable. And as CAD pattern recognition schemes become more advanced, these on-site adjustments are likely to become more complex. One hurdle to significantly improved CAD, then, is the added burden that more advanced pattern recognition schemes place on the sonar operator, to whom it falls to make the final system adjustments, by manually entering parameters that presume an ever greater working knowledge of image processing. This paper presents the concept of automatic and non-intrusive on-site training for CAD, in which the system gathers its training about targets and clutter at the start of a mission, simply by "eavesdropping" on the sonar operator's routine responses to the sonar image (selecting potential threats and passing over clutter as ordinarily expected of the operator). This automatic approach makes it possible to use more advanced pattern recognition, achieving better CAD performance without sacrificing mission flexibility or taxing the sonar operator.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document They should be selected so that no security classification is required Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included If possible keywords should be selected from a published thesaurus e g Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title)

mine hun`ting, sidescan sonar, pattern recognition, detection, expert system

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

Defence R&D Canada
is the national authority for providing
Science and Technology (S&T) leadership
in the advancement and maintenance
of Canada's defence capabilities.

R et D pour la défense Canada
est responsable, au niveau national, pour
les sciences et la technologie (S et T)
au service de l'avancement et du maintien des
capacités de défense du Canada.

#516765
CA020123

DEFENCE R&D DÉFENSE

www.drdc-rddc.dnd.ca